

# Almost Universal Hash Families are also Storage Enforcing

Mohammad Iftekhhar Husain      Steve Ko      Atri Rudra\*      Steve Uurtamo\*

Department of Computer Science and Engineering,  
University at Buffalo, SUNY,  
Buffalo, NY, 14214.  
{imhusain,steveko,atri,uurtamo}@buffalo.edu

## Abstract

We show that every almost universal hash function also has the *storage enforcement* property. Almost universal hash functions have found numerous applications and we show that this new storage enforcement property allows the application of almost universal hash functions in a wide range of remote verification tasks: (i) Proof of Secure Erasure (where we want to remotely erase and securely update the code of a compromised machine with memory-bounded adversary), (ii) Proof of Ownership (where a storage server wants to check if a client has the data it claims to have before giving access to deduplicated data) and (iii) Data possession (where the client wants to verify whether the remote storage server is storing its data). Specifically, storage enforcement guarantee in the classical data possession problem removes any practical incentive for the storage server to cheat the client by saving on storage space.

The proof of our result relies on a natural combination of Kolmogorov Complexity and List Decoding. To the best of our knowledge this is the first work that combines these two techniques. We believe the newly introduced storage enforcement property of almost universal hash functions will open promising avenues of exciting research under memory-bounded (bounded storage) adversary model.

**Keywords:** Kolmogorov Complexity, List Decoding, Almost Universal Hash family, Data Possession, Proof of Retrievability, Proof of Ownership, Reed-Solomon Codes, CRT codes.

---

\*Supported in part by NSF CAREER grant CCF-0844796

# 1 Introduction

Universal hashing was defined in the seminal paper of Carter and Wegman [6] and since then has been an integral part of both complexity and algorithms research. In this paper we consider the following well known generalization of universal hashing called  $\varepsilon$ -almost universal [28]. In particular, a family  $\mathcal{H} = \{H_1, \dots, H_n\}$  of hash functions (where  $H_i : [q]^k \rightarrow [q]$ )<sup>1</sup> is an  $\varepsilon$ -almost universal hash family if for every  $x \neq y \in [q]^k$ , we have

$$\Pr_{i \in [n]} [h_i(x) = h_i(y)] \leq \varepsilon.$$

( $1/q$ -almost universal hash families are universal hash families as defined in [6].)

In this paper, we show that any  $\varepsilon$ -almost universal hash family also has the “storage enforcement” property. In particular, we call a hash function family  $\mathcal{H} = \{H_1, \dots, H_n\}$  to be  $(\gamma, f(\cdot))$ -storage enforceable if the following holds. A *prover* claims to have  $x \in [q]^k$ . The prover is allowed to perform arbitrary computation on  $x$  and retain its output  $y \in [q]^*$  with it. Then the following holds for any  $x \in [q]^k$ : if the prover can compute<sup>2</sup>  $h_i(x)$  with probability at least  $\gamma$  (for a uniformly random  $i \in [n]$ ), then  $|y| \geq f(x)$ .

Our main result is the following:

**Theorem 1.** *If a computable hash function family  $\mathcal{H} = \{H_1, \dots, H_n\}$  is  $\varepsilon$ -almost universal, then it is also  $(\sqrt{\varepsilon}, f(\cdot))$ -storage enforceable, where  $f(x) \approx C(x)$ , where  $C(x)$  is the plain Kolmogorov complexity of  $x$ .*

The storage enforcement property is interesting in its own right and we present its applications in problems of proof of secure erasure, proof of ownership and data possession.

Before discussing the applications of storage enforcement, we present the techniques we use to prove Theorem 1, which to the best of our knowledge is the first instance of combining list decoding [9, 31] and Kolmogorov complexity [19, 20] naturally to give interesting results. Both Kolmogorov complexity (see the textbook [20]) and list decoding (see the survey by Sudan [29] and Guruswami’s thesis [14]) have found numerous applications in complexity theory. We believe that this combination merits more exploration and hope that this work will lead to a more systematic study.

**Our techniques.** We begin our discussion with the following related result. Let  $H : [q]^k \rightarrow [q]^n$  be an error-correcting code with good *list decodability* properties— see Section 2.1 for coding definitions. (We do not need an algorithmic guarantee, as just a combinatorial guarantee suffices.) Consider the natural hash family  $\mathcal{H} = \{h_1, \dots, h_n\}$ , where  $h_i(x) = H(x)_i$ . We now argue that  $\mathcal{H}$  has storage enforcement property. In particular, we quickly sketch why the prover cannot get away with storing a vector  $y$  such that  $|y|$  is smaller than  $C(x)$ , the plain Kolmogorov complexity of  $x$ , by some appropriately small additive factor. Since we are assuming that the prover uses an algorithm  $\mathcal{A}_x$  to compute its answer  $\mathcal{A}_x(\beta, y)$  to the random challenge  $\beta \in [n]$ , if the prover’s answer is correct with probability at least  $\gamma$ , then note that the vector  $(\mathcal{A}_x(\beta, y))_{\beta \in [n]}$  differs from  $H(x)$  in at most  $1 - \gamma$  fraction of positions. Thus, if  $H$  has good list decodability, then (using  $\mathcal{A}_x$ ) one can compute a list  $\{x_1, \dots, x_L\}$  that contains  $x$ . Finally, one can use  $\log L$  bits of advice (in addition to  $y$ ) to output  $x$ . This procedure then becomes a description for  $x$  and if  $|y|$  is sufficiently smaller than  $C(x)$ , then our description will have size  $< C(x)$ , which contradicts the definition of  $C(x)$ .

The above result then implies Theorem 1 because (1) Every  $\varepsilon$ -almost universal hash family defines a code  $H$  (in exactly the same way as we defined corresponding hash family  $\mathcal{H}$  from the code  $H$  in the

<sup>1</sup>Our results can also handle the case where each  $H_i : [q]^k \rightarrow [q_i]$  for potentially different  $q_i$  for each  $i \in [n]$ . For simplicity, we will concentrate on the case of  $q_i = q$  for every  $i \in [n]$ .

<sup>2</sup>We will assume the prover uses an algorithm  $\mathcal{A}_x$  to compute  $\mathcal{A}_x(i, y)$  as its version of  $h_i(x)$ .

paragraph above) with (relative) distance at least  $1 - \varepsilon$ ; (2) Every code  $H$  with relative distance  $1 - \varepsilon$  can be (combinatorially) list decoded from  $1 - \sqrt{\varepsilon}$  fraction of errors by the Johnson bound (cf. [14]).

From a more practical point of view, the Karp Rabin hash corresponds to  $H$  being so-called the ‘‘Chinese Remainder Theorem’’ (CRT) code and the polynomial hash corresponds to  $H$  being the Reed-Solomon code. Reed-Solomon and CRT codes have good list decodability, which implies that our protocols can be implemented using the classic Karp Rabin and polynomial hashes.

Next, we present three applications of hash functions that have storage enforcement.

**Proof of Secure Erasure.** Often, it is important to verify the internal state of a remote embedded device such as a sensor, actuator or computer peripheral with limited memory. This is to assure that it is running the intended code and not a malicious or arbitrary code. When any suspicious behavior of the remote device is detected, sometimes it becomes necessary to erase the entire memory content and then update it with legitimate code. Two crucial components of this process is to make sure that (i) the remote device has to store a random string at least as large as the memory size and (ii) efficiently verify that the device has indeed stored the random string remotely. A similar problem has been studied before, under the name of *Proof of Secure Erasure* [21] using cryptographic primitives which is, in most cases, too expensive for resource-constrained embedded devices. However, almost universal hash functions with the storage enforcement property applies naturally to provide an efficient proof of secure erasure scheme as outlined below.

The prover can use any storage enforceable hash family, along with a randomly-generated string  $x$ , (which will, with constant probability greater than  $1/2$ , have  $C(x) \geq (|x| - O(1))^3$ ), to force the remote device (verifier) to first store a string of length close to  $|x|$  (overwriting any malicious code as long as  $|x|$  is chosen large enough), second, to correctly answer a verification request (a failure to answer the request proves that the initial erasure has failed), and third (now that there is no room for a malicious code) to install the updated code.

Apart from efficiency issues, the cited work [21] also assumes that the remote device stores some portion of its trusted code in a fixed ROM. However, a universal hash function based scheme with storage enforcement is not strictly dependent on such an assumption. We do, however, (as does [21]) require that the remote machine not have network access to an external storage which is a very reasonable assumption in memory bounded adversary model.

**Proof of Ownership.** Consider the case when a client wants to upload its data  $x$  to a storage service (such as Dropbox) that uses ‘deduplication.’ In particular, if  $x$  is already stored on the server, then the server will ask the client not to upload  $x$  and give it ownership of  $x$ . To save on communication, the server asks the client to send it a hash  $h(x)$  and if it matches the hash of the stored  $x$  on the server, the client is issued an ownership of  $x$ . As identified by [16], this simple deduplication scheme can be abused for malicious purposes. For example, if a malicious user gets hold of the hashes of the files stored on the server (through server break-in or unintended leakage), it will get access to those files. The simple deduplication scheme can also be exploited as a unintended content distribution network. A malicious user can upload a potentially large file and share the hash of the file with accomplices. Now, all these accomplices can present the hash to storage server and get access to large file as if it is a content distribution network. A storage enforceable hash function can address such a situation. Specifically, if one has a hash family  $\{h_1, \dots, h_n\}$  that is storage enforceable, then the server can pick a random  $i \in [n]$  and ask the client to send it  $h_i(x)$ . If the client succeeds in doing so then, the server knows the client has a lot of information (close to  $C(x)$ ) related to  $x$  and not just a small fingerprint of  $x$ .

---

<sup>3</sup>The exact bound is very strong.  $C(x) \geq (|x| - r)$  for at least a  $(1 - 1/2^r)$ -fraction of the strings of length  $|x|$ .

Proof of ownership is studied formally in [16]. In addition to the scheme being computation intensive, the security guarantee in [16] is based on the min-entropy of the original file dealing with the probability distribution of input data blocks, which we believe (as also agreed by the authors of [16]) is not the best measure to enforce storage. We do not make any assumptions on what kind of data we are storing. Due to this, Kolmogorov complexity is a better option as it is defined for each specific block of data, which is at the core of our solution. Also, majority of the existing data possession schemes (as discussed in the next application) could not be used to address proof of ownership. Because, they require the prover (client) to store  $x$  in a modified form which is an infeasible assumption for this application.

**Data Possession.** This application considers the following problem: A string  $x \in [q]^k$  is held by the client/user, then transmitted to a remote server (or broken up into several pieces and transmitted to several remote servers). At some later point, after having sent  $x$ , the client would like to know that the remote server(s) is storing  $x$  correctly. The problem has been studied under the name of *proof of retrievability* (cf. [17, 5, 8]) and/or *data possession* (cf. [1, 2, 7, 10]). Given the greater prevalence of outsourcing of storage (e.g. in “cloud computing”), such a verification procedure is crucial for auditing purposes to make sure that the remote servers are following the terms of their agreement with the client.

In contrast with the existing schemes, our approach to the remote server adversarial model is rather practical. We assume that the storage provider is not fully malicious: e.g., it is reasonable to assume that Amazon will not try and mess with an individual’s storage in a malicious way. However, the server will not be completely honest either. We make the natural assumption that the servers’ main goal is to save on its storage space (since it practically helps the server to save on the cost of storage medium, networking and power): so if it can pass the verification protocol by storing substantially less than  $|x|$  amounts of data, then it will do so.<sup>4</sup> In our scheme, we will strive to make this impossible for the server to do without having a high probability of getting caught.

Also, before we move on to our scheme, we would like to make a point that does not seem to have been made explicitly before. Note that we cannot prevent a server from compressing their copy of the string, or otherwise representing it in any reversible way. (Indeed, the user would not care as long as the server is able to recreate  $x$ .) This means that a natural upper bound on the amount of storage we can force a server into is  $C(x)$ , the (plain) Kolmogorov complexity of  $x$ , which is the size of the smallest (algorithmic) description of  $x$ .

Our scheme to address the data possession problem is determined by a  $(\gamma, f(\cdot))$ -storage enforceable hash family  $\{h_1, \dots, h_n\}$ . In the pre-processing step, the client picks a random  $i \in [n]$ , stores  $(i, h_i(x))$  and ships off  $x$  to the server. During the verification step, the client send  $i$  and asks for  $h_i(x)$ . Upon receiving the reply  $z$  from the server, the client checks whether  $z = h_i(x)$ . Since the hash family was storage enforceable, if the server passes the verification protocol with probability  $\gamma > 0$ , then they provably have to store  $f(x)$  (which in our results will be  $C(x)$  up to a very small *additive* factor) bits of data. A cheating server is allowed to use arbitrarily powerful algorithms (as long as they terminate) while responding to challenges from the user.<sup>5</sup> Further, every honest server only needs to store  $x$ . In other words, unlike some existing results, our protocol does not require the server(s) to store a massaged version of  $x$ . In practice, this is important as we do not want our auditing protocol to interfere with “normal” read operations. However, unlike many of the existing results based upon cryptographic primitives, our protocol can only allow a number of audits that

---

<sup>4</sup>In other words, the server is willing to spend a lot of computational power to save on its space usage but is not maliciously trying to cheat the client.

<sup>5</sup>The server(s) can use different algorithms for different strings  $x$  but the algorithm cannot change across different challenges for the same  $x$ .

is proportional to the user’s local storage.<sup>6</sup> Further, as a demonstration of the versatility of our technique, it’ll turn out that our schemes can also provide a proof of retrievability (see e.g. Remark 1), but we do not claim any improvement over existing work on proofs of retrievability, e.g. [8]. In general, our single prover scheme is very similar to one of the schemes discussed in [8]. For example, one of the instantiations of our scheme with Reed-Solomon codes in fact resembles exactly one of their scheme. However, some crucial differences in technical contribution of the schemes are the following. First, even though both schemes use list decodability of codes, the schemes differ in how one “prunes” the output of list decoding. The scheme in [8] uses extra hash functions while our work uses Kolmogorov complexity. Second and perhaps more importantly, we believe that the main contribution of our work is to formally establish the storage enforcement properties for every almost universal hash function, and showing a range of applications of these storage enforceable hash functions.

Our techniques are very flexible. As mentioned earlier, our protocols also imply a proof of retrievability. As a further demonstration of this flexibility, we show that our results seamlessly generalize to the multiple server case. We show that we can allow for arbitrary collusion among the servers if we are only interested in checking if at least one sever is cheating (but we cannot identify it). If we want to identify at least one cheating server (and allow for servers to not respond to challenges), then we can handle a moderately limited collusion among servers.

Theorem 1 along with the above then implies that any almost universal hash family implies a valid verification protocol for the data possession problem. Our protocol with best parameters, which might need the user and honest server(s) to be exponential time algorithms, provably achieves the optimal local storage for the user and the optimal total communication between the user and the server(s). By picking an appropriate almost universal hash family, with slightly non-optimal storage and communication parameters, the user and the honest server(s) can work with single pass logarithmic space data stream algorithms.

## 2 Preliminaries

We now formally define different parameters of a verification protocol based on storage enforceable hash functions. For the secure erasure application, the administrator/user is the verifier and the remote embedded device is the prover. In the proof of ownership problem, the deduplication server is the verifier and the client is the prover. For the data possession problem, the client is the verifier and the storage server is the prover. For clarity, rest of the discussion will focus on the data possession application to accomodate both single prover and multiple prover cases.

**Basic Verification Scheme.** We use  $\mathbb{F}_q$  to denote the finite field over  $q$  elements. We also use  $[n]$  to denote the set  $\{1, 2, \dots, n\}$ . Given any string  $x \in [q]^*$ , we use  $|x|$  to denote the length of  $x$  in bits. Additionally, all logarithms will be base 2 unless otherwise specified. We use  $U$  to denote the *verifier*.

For the data possession problem, we define the verification protocol in terms of multiple prover case first and then discuss single prover as a specific case for ease of dicussion. We assume that  $U$  wants to store its data  $x \in [q]^k$  among  $s$  service providers (or provers)  $P_1, \dots, P_s$ . For the proof of ownership problem, the client (or the prover)  $P_1$  has the data  $x$ . In the preprocessing step, the provers  $\mathcal{P} = \{P_1, \dots, P_s\}$  get  $x$

---

<sup>6</sup>An advantage of proving security under cryptographic assumptions is that one can leverage existing theorems to prove additional security properties of the verification protocol. We do not claim any additional security guarantees other than the ability of being able to force servers to store close to  $C(x)$  amounts of data.

divided up equally among the  $s$  provers – we will denote the chunk for prover  $i \in [s]$  as  $x_i \in [q]^{n/s}$ .<sup>7</sup> (For the data possession problem, we can assume that  $U$  does this division and sends  $x_i$  to  $P_i$ .) Each prover is then allowed to apply any computable function to its chunk and to store a string  $y_i \in [q]^*$ . Ideally, we would like  $y_i = x_i$ . However, since the provers can compress  $x$ , we would at the very least like to force  $|y_i|$  to be as close to  $C(x_i)$  as possible. For notational convenience, for any subset  $T \subseteq [s]$ , we denote  $y_T$  ( $x_T$  resp.) to be the concatenation of the strings  $\{y_i\}_{i \in T}$  ( $\{x_i\}_{i \in T}$  resp.).

To enforce the conditions above, we design a protocol. We will be primarily concerned with the amount of storage at the verifier side and the amount of communication and want to minimize both simultaneously while giving good verification properties. The following definition captures these notions. (We also allow for the provers to collude among each other.)

**Definition 1.** Let  $s, c, m \geq 1$  and  $0 \leq r \leq s$  be integers,  $0 \leq \rho \leq 1$  be a real and  $f : [q]^* \rightarrow \mathbb{R}_{\geq 0}$  be a function. Then an  $(s, r)$ -party verification protocol with resource bound  $(c, m)$  and verification guarantee  $(\rho, f)$  is a randomized protocol with the following guarantee. For any string  $x \in [q]^k$ ,  $U$  stores at most  $m$  bits and communicates at most  $c$  bits with the  $s$  provers. At the end, the protocol either outputs a 1 or a 0. Finally, the following is true for any  $T \subseteq [s]$  with  $|T| \leq r$ : If the protocol outputs a 1 with probability at least  $\rho$ , then assuming that every prover  $i \in [s] \setminus T$  followed the protocol and that every prover in  $T$  possibly colluded with one another, we have  $|y_T| \geq f(x_T)$ .

We will denote a  $(1, 1)$ -party verification protocol as a one-party verification protocol. (Note that in this case, the single prover is allowed to behave arbitrarily.)

All of our protocols will have the following structure: we first pick a hash family. The protocol will pick random hash(es) and store the corresponding hash values for  $x$  (along with the indices of the hash functions) during the pre-processing step. During the verification step,  $U$  sends the indices to the hashes as challenges to the  $s$  provers. Throughout this paper, we will assume that each prover  $i$  has a computable algorithm  $\mathcal{A}_{x,i}$  such that on challenge  $\beta$  it returns an answer  $\mathcal{A}_{x,i}(\beta, y_i)$  to  $U$ . The protocol then outputs 1 or 0 by applying a (simple) boolean function on the answers and the stored hash values.

In particular, for the one-party verification protocol, we have the following simple result, which we prove in the appendix.

**Proposition 2.** Let  $\mathcal{H} = \{h_1, \dots, h_n\}$  be a hash family where  $h_i : [q]^k \rightarrow [q]$  that is  $(\gamma, f)$ -storage enforceable. Then there exists a one party verification protocol with resource bound  $(\log q + \log n, \log q + \log n)$  and performance guarantee  $(\gamma, f)$ .

## 2.1 Coding Basics

We begin with some basic coding definitions. An error-correcting code  $H$  with dimension  $k \geq 1$  and block length  $n \geq k$  over an alphabet of size  $q$  is any function  $H : [q]^k \rightarrow [q]^n$ . A linear code  $H$  is any error-correcting code that is a linear function, in which case we correspond  $[q]$  with  $\mathbb{F}_q$ . A message of a code  $H$  is any element in the domain of  $H$ . A codeword in a code  $H$  is any element in the range of  $H$ . The Hamming distance  $\Delta(x, y)$  of two same-length strings is the number of symbols in which they differ. The relative distance  $\delta$  of a code is  $\min_{x \neq y} \frac{\Delta(x, y)}{n}$ , where  $x$  and  $y$  are any two different codewords in the code.

The following connection between almost universal and codes is well known. The straightforward proof is in the appendix.

<sup>7</sup>We will assume that  $s$  divides  $n$ . In our results for the case when  $H$  is a linear code, we do not need the  $x_i$ 's to have the same size, only that  $x$  can be partitioned into  $x_1, \dots, x_s$ . We will ignore this possibility for the rest of the paper.



**Proposition 3** ([4]). Let  $\mathcal{H} = \{h_1, \dots, h_n\}$  be an  $\varepsilon$ -almost universal hash family with  $h_i : [q]^k \rightarrow [q]$ . Now consider the related code  $q$ -ary code  $H$  with block length  $n$  and dimension  $k$  where  $H(x) = (h_i(x))_{i \in [n]}$ . Then  $H$  has relative distance at least  $1 - \varepsilon$ .

**Definition 2.** A  $(\rho, L)$  list-decodable code is any error-correcting code such that for every vector  $e$  in the codomain of  $H$ , the number of codewords that are Hamming distance  $\rho n$  or less from  $e$  is always  $L$  or fewer.

**Johnson Bound.** We now state a general combinatorial result for list decoding codes with large distance, which will be useful. The result below allows for a sort of non-standard definition of codes, where a codeword is a vector in  $\prod_{i=1}^n [q_i]$ , where the  $q_i$ 's can be distinct.<sup>8</sup> (So far we have looked only at the case where  $q_i = q$  for  $i \in [n]$ .) The notion of Hamming distance still remains the same, i.e. the number of positions that two vectors differ in. (The syntactic definitions of the distance of a code and the list decodability of a code remain the same.) We will need the following result:

**Theorem 4** ([14]). Let  $C$  be a code with block length  $n$  and distance  $d$  where the  $i$ th symbol in a codeword comes from  $[q_i]$ . Then the code is  $\left(1 - \sqrt{1 - \frac{d}{n}}, 2 \sum_{i=1}^n q_i\right)$ -list decodable.

Theorem 4 and Proposition 3 implies the following:

**Corollary 5.** Let  $\mathcal{H} = \{h_1, \dots, h_n\}$  be an  $\varepsilon$ -almost universal hash family with  $h_i : [q]^k \rightarrow [q]$ . Now consider the related code  $q$ -ary code  $H$  where  $H(x) = (h_i(x))_{i \in [n]}$ . Then  $H$  is  $(1 - \sqrt{\varepsilon}, 2qn)$ -list decodable.

For the purposes of this paper, we only consider codes  $H$  that are members of a family of codes  $\bar{\mathcal{H}}$ , any one of which can be indexed by  $(k, n, q, \rho n)$ .

### Plain Kolmogorov Complexity.

**Definition 3.** The plain Kolmogorov Complexity  $C(x)$  of a string  $x$  is the minimum sum of sizes of a compressed representation of  $x$ , along with its decompression algorithm  $D$ , and a reference universal Turing machine that runs the decompression algorithm.

Because the reference universal Turing machine size is constant, it is useful to think of  $C(x)$  as simply measuring the amount of inherent (i.e. incompressible) information in a string  $x$ .

## 3 Single Prover Storage Enforcement

We begin by presenting our main result for the case of one prover ( $s = 1$ ) to illustrate the combination of list decoding and Kolmogorov complexity. In the subsequent section, we will generalize our result to the multiple prover case (for the data possession problem).

**Theorem 6.** For every computable error-correcting code  $H : [q]^k \rightarrow [q]^n$  that is  $(\rho, L)$  list-decodable, define the hash family  $\mathcal{H} = \{h_1, \dots, h_n\}$  as  $h_i(x) = H(x)_i$ . Then  $\mathcal{H}$  is a  $(1 - \rho, f)$ -storage enforceable, where for any  $x \in [q]^k$ ,  $f(x) = C(x) - \log(qLn^3) - 2 \log \log(qn) - c_0$ , for some fixed constant  $c_0$ .<sup>9</sup>

<sup>8</sup>We're overloading the product operator  $\prod$  here to mean the iterated Cartesian product.

<sup>9</sup>The contribution from the encoding of the constants in this theorem to  $c_0$  is 2. For most codes, we can take  $c_0$  to be less than a few thousand. What is important is that the contribution from the encoding is independent of the rest of the constants in the theorem. Although we do not explicitly say so in the body of the theorem statements, this important fact is true for the rest of the results in this paper as well.

*Proof.* We assume that the prover, upon receiving  $x$ , saves a string  $y \in [q]^*$ . The prover is allowed to use any computable function to obtain  $y$  from  $x$ . Further, we assume that the prover, upon receiving the (random) challenge  $\beta \in [n]$ , uses a computable function  $\mathcal{A}_x : [n] \times [q]^* \rightarrow [q]$  to compute  $a = \mathcal{A}_x(\beta, y)$ .

We prove the storage enforceable property by contradiction. Assume that  $|y| < f(x) \stackrel{\text{def}}{=} C(x) - \log(qLn^3) - 2 \log \log(qn) - c_0$  and yet  $\mathcal{A}_x(\beta, y) = h_\beta(x) = H(x)_\beta$  with probability at least  $1 - \rho$  (over the choice of  $\beta$ ). Define  $z = (\mathcal{A}_x(\beta, y))_{\beta \in [n]}$ . Note that by the claim on the probability,  $\Delta(z, H(x)) \leq \rho n$ . We will use this and the list decodability of the code  $H$  to prove that there is an algorithm with description size  $< C(x)$  to describe  $x$ , which is a contradiction. To see this, consider the following algorithm that uses  $y$  and an advice string  $v \in \{0, 1\}^{|L|}$ :

1. Compute a description of  $H$  from  $n, k, \rho n$  and  $q$ .
2. Compute  $z = (\mathcal{A}_x(\beta, y))_{\beta \in [n]}$ .
3. By cycling through all  $x \in [q]^k$ , retain the set  $\mathcal{L} \subseteq [q]^k$  such that for every  $u \in \mathcal{L}$ ,  $\Delta(H(u), z) \leq \rho n$ .
4. Output the  $v$ th string from  $\mathcal{L}$ .

Note that since  $H$  is  $(\rho, L)$ -list decodable, there exists an advice string  $v$  such that the algorithm above outputs  $x$ . Further, since  $H$  is computable, there is an algorithm  $\mathcal{E}$  that can compute a description of  $H$  from  $n, k, \rho n$  and  $q$ . (Note that using this description, we can generate any codeword  $H(u)$  in step 3.) Thus, we have description of  $x$  of size  $|y| + |v| + |\mathcal{A}_x| + |\mathcal{E}| + (3 \log n + \log q + 2 \log \log n + 2 \log \log q + 2)$  (where the last term is for encoding the different parameters<sup>10</sup>), which means if  $|y| < C(x) - |v| - |\mathcal{A}_x| - |\mathcal{E}| - (3 \log n + \log q + 2 \log \log n + 2 \log \log q + 2) = f(x)$ , then we have a description of  $x$  of size  $< C(x)$ , which is a contradiction. □

The one unsatisfactory aspect of the result above is that if  $H$  is not polynomial time computable, then Step 2 in the pre-processing step for  $U$  is not efficient. Similarly, if the server is not cheating (and e.g. stores  $y = x$ ), then it cannot also compute the correct answer efficiently. We will come back to these issues when we instantiate  $H$  by an explicit code such as Reed-Solomon.

**Remark 1.** If  $H$  has relative distance  $\delta$ , then note that  $\mathcal{H}$  is  $(1 - \delta/2 - \epsilon, C(x) - \log(qn^3) - 2 \log \log(qn) - c_0)$ -storage enforceable for some fixed constant  $c_0$ . Further,  $y$  has enough information for the prover to compute  $x$  back from it. (It can use the same algorithm to compute  $x$  from  $y$  detailed above, except it does not need the advice string  $v$ , as in Step 3 we will have  $\mathcal{L} = \{x\}$ .) For the more general case when  $H$  is  $(\rho, L)$ -list decodable and  $\mathcal{H}$  is  $(1 - \rho, C(x) - \log qLn^3 - 2 \log \log qn - c_0)$ -storage enforceable, then  $y$  has enough information for the prover to compute a list  $\mathcal{L} \supseteq \{x\}$  with  $|\mathcal{L}| \leq L$ . The verifier, if given access to  $\mathcal{L}$ , can use its randomly stored  $h_\beta(x)$  to pick  $x$  out of  $\mathcal{L}$  with probability at least  $1 - \delta L$ .

Theorem 6 along with Proposition 2 implies the following result:

**Theorem 7.** For every computable error-correcting code  $H : [q]^k \rightarrow [q]^n$  that is  $(\rho, L)$  list-decodable, there exists a one-party verification protocol with resource bound  $(\log n + \log q, \log n + \log q)$  and verification guarantee  $(1 - \rho, f)$ , where for any  $x \in [q]^k$ ,  $f(x) = C(x) - \log(qLn^3) - 2 \log \log(qn) - c_0$ , for some fixed constant  $c_0$ .

---

<sup>10</sup>We use a simple self-delimiting encoding of  $q$  and  $n$ , followed immediately by  $k$  and  $\rho n$  in binary, with the remaining bits used for  $v$ . A simple self-delimiting encoding for a positive integer  $u$  is the concatenation of:  $(\lceil \log(|u|) \rceil)$  in unary, 0,  $|u|$  in binary,  $u$  in binary. We omit the description of this encoding in later proofs.



## 4 Multiple Prover Storage Enforcement

In this section, we show how to extend the result from Theorem 7 to the multiple server case. In the first two sub-sections, we will implicitly assume the following: (i) We are primarily interested in whether some server was cheating and not in identifying the cheater(s) and (ii) We assume that all servers always reply back (possibly with an incorrect answer).

**Trivial Solution.** We begin with the following direct generalization of Theorem 7 to the multiple server case: essentially run  $s$  independent copies of the protocol from Theorem 7.

**Theorem 8.** *For every computable error-correcting code  $H : [q]^{k/s} \rightarrow [q]^n$  that is  $(\rho, L)$  list-decodable, there exists an  $(s, s)$ -party verification protocol with resource bound  $(\log n + s \log q, s(\log n + \log q))$  and verification guarantee  $(1 - \rho, f)$ , where for any  $x \in [q]^k$ ,  $f(x) = C(x) - s - \log(s^2 q L^s n^4) - 2 \log \log(qn) - c_0$ , for some fixed positive integer  $c_0$ .*

**Multiple Parties, One Hash.** One somewhat unsatisfactory aspect of Theorem 8 is that the storage needed by  $U$  goes up a factor of  $s$  from that in Theorem 7. Next we show that if the code  $H$  is linear (and list decodable) then we can get a similar guarantee as that of Theorem 8 except that the storage usage of  $U$  remains the same as that in Theorem 7.

**Theorem 9.** *For every computable linear error-correcting code  $H : \mathbb{F}_q^k \rightarrow \mathbb{F}_q^n$  that is  $(\rho, L)$  list-decodable, there exists an  $(s, s)$ -party verification protocol with resource bound  $(\log n + \log q, s(\log n + \log q))$  and verification guarantee  $(1 - \rho, f)$ , where for any  $x \in \mathbb{F}_q^k$ ,  $f(x) = C(x) - s - \log(s^2 q L n^4) - 2 \log \log(qn) - c_0$ , for some fixed positive integer  $c_0$ .*

**Catching the Cheaters and Handling Unresponsive Servers.** The protocol in Theorem 8 checks if the answer from each server is the same as corresponding stored hash. This implies that the protocol can easily handle the case when some server does not reply back at all. Additionally, if the protocol outputs a 0 then it knows that at least one of the servers in the colluding set is cheating. (It does not necessarily identify the exact set  $T$ .<sup>11</sup>)

However, the protocol in Theorem 9 cannot identify the cheater(s) and needs all the servers to always reply back. Next, using Reed-Solomon codes, at the cost of higher user storage and a stricter bound on the number of colluding servers, we show how to get rid of these shortcomings.

Recall that a Reed-Solomon code  $RS : \mathbb{F}_q^m \rightarrow \mathbb{F}_q^n$  can be represented as a systematic code (i.e. the first  $k$  symbols in any codeword is exactly the corresponding message) and can correct  $r$  errors and  $e$  erasures as long as  $2r + e \leq n - m$ . Further, one can correct from  $r$  errors and  $e$  erasures in  $O(3)$  time. The main idea in the following result is to follow the protocol of Theorem 8 but instead of storing all the  $s$  hashes,  $U$  only stores the parity symbols in the corresponding Reed-Solomon codeword.

**Theorem 10.** *For every computable linear error-correcting code  $H : \mathbb{F}_q^k \rightarrow \mathbb{F}_q^n$  that is  $(\rho, L)$  list-decodable, assuming at most  $e$  servers will not reply back to a challenge, there exists an  $(r, s)$ -party verification protocol with resource bound  $(\log n + (2r + e) \cdot \log q, s(\log n + \log q))$  and verification guarantee  $(1 - \rho, f)$ , where for any  $x \in \mathbb{F}_q^k$ ,  $f(x) = C(x) - s - \log(s^2 q L n^4) - 2 \log \log(qn) - c_0$ , for some fixed positive integer  $c_0$ .*

---

<sup>11</sup>We assume that identifying at least one server in the colluding set is motivation enough for servers not to collude.

## 5 Corollaries

We now present specific instantiations of list decodable codes  $H$  to obtain corollaries of our main results.

**Optimal Storage Enforcement.** We begin with the following observation: If the reply from a prover comes from a domain of size  $q$ , then one cannot hope to have a  $(\delta, f)$ -storage enforceable hash family for any  $\delta \leq 1/q$  for any non-trivial  $f$ . This is because the prover can always return a random value and guess any  $h_i(x)$  with probability  $1/q$ .

Next, we show that we can get  $\delta$  to be arbitrarily close to  $1/q$  while still obtaining  $f(x)$  to be very close to  $C(x)$ . We start off with the following result due to Zyablov and Pinsker:

**Theorem 11** ([32]). *Let  $q \geq 2$  and let  $0 < \rho < 1 - 1/q$ . There exists a  $(\rho, L)$ -list decodable code with rate  $1 - H_q(\rho) - 1/L$ .*

It is known that for  $\varepsilon < 1/q$ ,  $H_q(1 - 1/q - \varepsilon) \leq 1 - C_q \varepsilon^2$ , where  $C_q = q/(4 \ln q)$  [24, Chap. 2]. This implies that there exists a code  $H : \mathbb{F}_q^k \rightarrow \mathbb{F}_q^n$ , with  $n \leq \frac{k}{(C_q - 1)\varepsilon^2} \leq 8k \ln q / (q\varepsilon^2)$ , which is  $(1 - 1/q - \varepsilon, 1/\varepsilon^2)$ -list decodable. Note that the above implies that one can deterministically compute a uniquely-determined such code by iterating over all possible codes with dimension  $k$  and block length  $n$  and outputting the lexicographically least such one that is  $(1 - 1/q - \varepsilon, L)$ -list decodable with the smallest discovered value of  $L$ . Applying this to Theorem 6 implies the following optimal result:

**Corollary 12.** *For every  $\varepsilon < 1/q$  and integer  $s \geq 1$ , there exists a  $(1/q + \varepsilon, f)$ -storage enforceable hash family, where for any  $x \in [q]^k$ ,  $f(x) = C(x) - s - \log s^2 k^4 + \log \varepsilon^{2s+8} - \log \log q^6 k^2 + \log \log \varepsilon^4 - \log \log \log q^3 - c_0$  for some fixed positive integer  $c_0$ .*

Some of our results need  $H$  to be linear. To this end, we will need the following result due to Guruswami et al.<sup>12</sup>

**Theorem 13** ([15]). *Let  $q \geq 2$  be a prime power and let  $0 < \rho < 1 - 1/q$ . Then a random linear code of rate  $1 - H_q(\rho) - \varepsilon$  is  $(\rho, C_{\rho,q}/\varepsilon)$ -list decodable for some term  $C_{\rho,q}$  that just depends on  $\rho$  and  $q$ .*

As a Corollary the above implies (along with the arguments used earlier in this section) that there exists a linear code  $H : \mathbb{F}_q^k \rightarrow \mathbb{F}_q^n$  with  $n \leq 8k \ln q / (q\varepsilon^2)$  that is  $(1 - 1/q - \varepsilon, C'_{\varepsilon,q}/\varepsilon^2)$ -list decodable (where  $C'_{\varepsilon,q} \stackrel{\text{def}}{=} C_{1-1/q-\varepsilon,q}$ ). Applying this to Theorem 10 gives us the following

**Corollary 14.** *For every  $\varepsilon < 1/q$ , integer  $s \geq 1$ , and  $r, e \leq s$ , assuming at most  $e$  servers do not reply back to a challenge, there exists an  $(r, s)$ -party verification protocol with resource bound  $(\log k q^{2r+e-1} - \log \varepsilon^2 + \log \log q^{3/2} + 3, s(\log k - \log \varepsilon^2 + \log \log q^{3/2} + 3))$  and verification guarantee  $(1/q + \varepsilon, f)$ , where for any  $x \in [q]^k$ ,  $f(x) = C(x) - s - \log s^2 C'_{\varepsilon,q} k^4 + \log q^3 \varepsilon^{10} - \log \log q^6 k^2 + \log \log \varepsilon^2 - \log \log \log q^2 - c_0$  for some fixed positive integer  $c_0$ .*

**Proof of Theorem 1.** Corollary 5 and Theorem 6 proves Theorem 1 with  $f(x) = C(x) - O(\log q + \log n)$ .

---

<sup>12</sup>The corresponding result for general codes has been known for more than thirty years.

**Practical Storage Enforcement.** All of our results so far have used (merely) computable codes  $H$ , which are not that useful in practice. What we really want in practice is to use codes  $H$  that lead to an efficient implementation of the protocol. At the very least, all the honest parties in the verification protocol should not have to use more than polynomial time to perform the required computation. An even more desirable property would be for honest parties to be able to do their computation in a one pass, logspace, data stream fashion. In this section, we'll see one example of each. Further, it turns out that the resulting hash functions are classical ones that are also used in practice. In particular, using the Karp-Rabin hash and the polynomial hash we get the following results. (More details are in the appendix.)

**Corollary 15.** *For every  $\varepsilon > 0$ , there exists an  $(s, s)$ -party verification protocol with resource bound*

$$((s+1)\log k + s - \log \varepsilon^4 + s \log \log(k/\varepsilon^2), s \log k - s \log \varepsilon^2 + s + s \log \log(k/\varepsilon^2))$$

*with verification guarantee  $(\varepsilon, f)$ , where for every  $x \in \{0, 1, \dots, \prod_{i=1}^k p_i - 1\}$ ,*

$$f(x) = C(x) - c_0(s(\log(k/\varepsilon) - \log \log(k/\varepsilon) - 1)) - c_1 \log \log \log(k/\varepsilon) - c_2$$

*for some fixed positive integers  $c_0, c_1$  and  $c_2$ . Further, all honest parties can do their computation in  $\text{poly}(n)$  time.*

**Corollary 16.** *For every  $\varepsilon > 0$ ,*

- (i) *There exists an  $(s, s)$ -party verification protocol with resource bound  $((s+1)(\log k + 2\log(1/\varepsilon) + 1), 2s(\log k + 2\log(1/\varepsilon) + 1))$  and verification guarantee  $(\varepsilon, f)$ , where for any  $x \in \mathbb{F}_q^k$ ,  $f(x) = C(x) - O(s(\log k + \log(1/\varepsilon)))$ .*
- (ii) *Assuming at most  $e$  servers do not respond to challenges, there exists an  $(r, s)$ -party verification protocol with resource bound  $((2r + e + 1)(\log k + 2\log(1/\varepsilon) + 1), 2s(\log k + 2\log(1/\varepsilon) + 1))$  and verification guarantee  $(\varepsilon, f)$ , where for any  $x \in \mathbb{F}_q^k$ ,  $f(x) = C(x) - O(s + \log k + \log(1/\varepsilon))$ .*

*Further, in both the protocols, honest parties can implement their required computation with a one pass,  $O(\log k + \log(1/\varepsilon))$  space (in bits) and  $\tilde{O}(\log k + \log(1/\varepsilon))$  update time data stream algorithm.*

## Acknowledgments

We thank Dick Lipton for pointing out the application of our protocol to the OS updating problem (a special case of proof of secure erasure) and for kindly allowing us to use his observation. We also thank Ram Sridhar for helpful discussions.

## References

- [1] G. Ateniese, R. Burns, R. Curtmola, J. Herring, L. Kissner, Z. Peterson, and D. Song. Provable data possession at untrusted stores. In *Proceedings of the 14th ACM conference on Computer and communications security*, CCS '07, pages 598–609, New York, NY, USA, 2007. ACM.
- [2] G. Ateniese, R. Di Pietro, L. V. Mancini, and G. Tsudik. Scalable and efficient provable data possession. In *Proceedings of the 4th international conference on Security and privacy in communication networks*, SecureComm '08, pages 9:1–9:10, New York, NY, USA, 2008. ACM.

- [3] E. Bach and J. Shallit. *Algorithmic number theory, volume 1: efficient algorithms*. MIT Press, Cambridge, Massachusetts, 1996.
- [4] J. Bierbrauer, T. Johansson, G. Kabatianskii, and B. J. M. Smeets. On families of hash functions via geometric codes and concatenation. In *CRYPTO*, pages 331–342, 1993.
- [5] K. D. Bowers, A. Juels, and A. Oprea. Proofs of retrievability: theory and implementation. In *Proceedings of the 2009 ACM workshop on Cloud computing security, CCSW '09*, pages 43–54, New York, NY, USA, 2009. ACM.
- [6] L. Carter and M. N. Wegman. Universal classes of hash functions. *J. Comput. Syst. Sci.*, 18(2):143–154, 1979.
- [7] R. Curtmola, O. Khan, R. Burns, and G. Ateniese. Mr-pdp: Multiple-replica provable data possession. *Distributed Computing Systems, International Conference on*, 0:411–420, 2008.
- [8] Y. Dodis, S. P. Vadhan, and D. Wichs. Proofs of retrievability via hardness amplification. In *Proceedings of the 6th Theory of Cryptography Conference (TCC)*, pages 109–127, 2009.
- [9] P. Elias. List Decoding for Noisy Channels. *Wescon Convention Record, Institute of Radio Engineers (now IEEE)*, Part 2:94–104, 1957.
- [10] C. Erway, A. Küpçü, C. Papamanthou, and R. Tamassia. Dynamic provable data possession. In *Proceedings of the 16th ACM conference on Computer and communications security, CCS '09*, pages 213–222, New York, NY, USA, 2009. ACM.
- [11] E. L. Gazzoni, D. Luiz, G. Filho, P. Sérgio, and L. M. Barreto. Demonstrating data possession and uncheatable data transfer. *Cryptology ePrint Archive*, Report 2006/150, 2006.
- [12] P. Golle, S. Jarecki, and I. Mironov. Cryptographic primitives enforcing communication and storage complexity. In *Proceedings of the 6th international conference on Financial cryptography, FC'02*, pages 120–135, Berlin, Heidelberg, 2003. Springer-Verlag.
- [13] G. R. Goodson, J. J. Wylie, G. R. Ganger, and M. K. Reiter. Efficient byzantine-tolerant erasure-coded storage. In *Proceedings of the 2004 International Conference on Dependable Systems and Networks*, pages 135–, Washington, DC, USA, 2004. IEEE Computer Society.
- [14] V. Guruswami. *List decoding of error-correcting codes*. Number 3282 in Lecture Notes in Computer Science. Springer, 2004.
- [15] V. Guruswami, J. Håstad, and S. Kopparty. On the list-decodability of random linear codes. *IEEE Transactions on Information Theory*, 57(2):718–725, 2011.
- [16] S. Halevi, D. Harnik, B. Pinkas, and A. Shulman-Peleg. Proofs of ownership in remote storage systems. *Cryptology ePrint Archive*, Report 2011/207, 2011.
- [17] A. Juels and B. S. K. Jr. Pors: proofs of retrievability for large files. In P. Ning, S. D. C. di Vimercati, and P. F. Syverson, editors, *ACM Conference on Computer and Communications Security*, pages 584–597. ACM, 2007.

- [18] R. M. Karp and M. O. Rabin. Efficient randomized pattern-matching algorithms. *IBM J. Res. Dev.*, 31:249–260, March 1987.
- [19] A. N. Kolmogorov. Three Approaches to the Quantitative Definition of Information. *Problems of Information Transmission*, 1(1):1–7, 1965.
- [20] M. Li and P. M. B. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Graduate Texts in Computer Science. Springer, New York, NY, USA, third edition, 2008.
- [21] D. Perito and G. Tsudik. Secure code update for embedded devices via proofs of secure erasure. In *Proceedings of the 15th European conference on Research in computer security, ESORICS’10*, pages 643–662, Berlin, Heidelberg, 2010. Springer-Verlag.
- [22] I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics*, 8(2):300–304, 1960.
- [23] J. B. Rosser and L. Schoenfeld. Sharper bounds for the chebyshev functions  $\theta(x)$  and  $\psi(x)$ . *Mathematics of Computation*, 29(129):pp. 243–269, 1975.
- [24] A. Rudra. *List Decoding and Property Testing of Error Correcting Codes*. PhD thesis, University of Washington, 2007.
- [25] T. Schwarz and E. L. Miller. Store, forget, and check: Using algebraic signatures to check remotely administered storage. In *Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS ’06)*, july 2006.
- [26] H. Shacham and B. Waters. Compact proofs of retrievability. In J. Pieprzyk, editor, *Advances in Cryptology - ASIACRYPT 2008*, volume 5350 of *Lecture Notes in Computer Science*, pages 90–107. Springer Berlin / Heidelberg, 2008.
- [27] M. A. Shah, R. Swaminathan, and M. Baker. Privacy-preserving audit and extraction of digital contents. Technical Report HPL-2008-32R1, HP Labs, 2008.
- [28] D. R. Stinson. Universal hashing and authentication codes. In *CRYPTO*, pages 74–85, 1991.
- [29] M. Sudan. List decoding: Algorithms and applications. *SIGACT News*, 31:16–27, 2000.
- [30] Q. Wang, C. Wang, J. Li, K. Ren, and W. Lou. Enabling public verifiability and data dynamics for storage security in cloud computing. In M. Backes and P. Ning, editors, *Computer Security ESORICS 2009*, volume 5789 of *Lecture Notes in Computer Science*, pages 355–370. Springer Berlin / Heidelberg, 2009.
- [31] J. M. Wozencraft. List decoding. *Quarterly Progress Report, Research Laboratory of Electronics, MIT*, 48:90–95, 1958.
- [32] V. V. Zyablov and M. S. Pinsker. List cascade decoding. *Problems of Information Transmission*, 17(4):29–34, 1981 (in Russian); pp. 236–240 (in English), 1982.



## A Related Works

Existing approaches for data possession verification at remote storage can be broadly classified into two categories: Crypto-based and Coding based. *Crypto-based approaches* rely on symmetric and asymmetric cryptographic primitives for proof of data possession. Ateniese et al. [1] defined the proof of data possession (PDP) model which uses public key homomorphic tags for verification of stored files. It can also support public verifiability with a slight modification of the original protocol by adding extra communication cost. In subsequent work, Ateniese et al. [2] proposed a symmetric crypto-based variation (SEP) which is computationally efficient compared to the original PDP but lacks public verifiability. Also, both of these protocols considered the scenario with files stored on a single server, and do not discuss erasure tolerance. However, Curtmola et al. [7] extended PDP to a multiple-server scenario by introducing multiple identical replicas of the original data. Among other notable constructions of PDP, Gazzoni et al. [11] proposed a scheme (DDP) that relied on an RSA-based hash (exponentiating the whole file), and Shah et al. [27] proposed a symmetric encryption based storage audit protocol. Recent extensions of crypto-based PDP schemes by Wang et al. (EPV) [30] and Erway et al. [10] mainly focus on supporting data dynamics in addition to existing capabilities. Golle et al. [12] had proposed a cryptographic primitive called storage enforcing commitment (SEC) which probabilistically guarantees that the server is using storage whose size is equal to the size of the original data to correctly answer the data possession queries. In general, the drawbacks of the aforementioned protocols are: (a) being computation intensive due to the usage of expensive cryptographic primitives and (b) since each verification checks a random fragment of the data, a small fraction of data corruption might go undetected and hence they do not guarantee the retrievability of the original data. *Coding-based approaches*, on the other hand, have relied on special properties of linear codes such as the Reed-Solomon (RS) [22] code. The key insight is that encoding the data imposes certain algebraic constraints on it which can be used to devise an efficient fingerprinting scheme for data verification. Earlier schemes proposed by Schwarz et al. (SFC) [25] and Goodson et al. [13] are based on this and are primarily focused on the construction of fingerprinting functions and categorically fall under distributed protocols for file integrity checking. Later, Juels and Kaliski [17] proposed a construction of a proof of retrievability (POR) which guarantees that if the server passes the verification of data possession, the original data is retrievable with high probability. While the scheme by Juels [17] supported a limited number of verifications, the theoretical POR construction by Shacham and Waters [26] extended it to unlimited verification and public verifiability by integrating cryptographic primitives. Subsequently, Dodis et al. [8] provided theoretical studies on different variants of existing POR schemes and Bowers et al. [5] considered POR protocols of practical interest [17, 26] and showed how to tune parameters to achieve different performance goals. However, these POR schemes only consider the single server scenario and have no construction of a retrievability and storage enforcement guarantee in a distributed storage scenario.

## B Omitted Proofs

### B.1 Proof of Proposition 2

*Proof.* We begin by specifying the protocol. In the preprocessing step, the verifier  $U$  does the following on input  $x \in [q]^k$ :

1. Generate a random  $\beta \in [n]$ .
2. Store  $(\beta, \alpha = h_\beta(x))$  (and in the data possession problem, send  $x$  to the prover).

The prover, upon receiving  $x$ , saves a string  $y \in [q]^*$ . The prover is allowed to use any computable function to obtain  $y$  from  $x$ .

During the verification phase,  $U$  does the following:

1. It sends  $\beta$  to the prover.
2. It receives  $a \in [q]$  from the prover. ( $a$  is supposed to be  $h_\beta(x)$ .)
3. It outputs 1 (i.e. prover did not “cheat”) if  $a = \alpha$ , else it outputs a 0.

The resource bound follows from the specification of the protocol and the performance guarantee follows from the fact that  $\mathcal{H}$  is  $(\gamma, f)$ -storage enforceable.  $\square$

## B.2 Proof of Proposition 3

*Proof.* This proof is straightforward. For  $x \neq y \in [q]^k$ , since  $\mathcal{H}$  is  $\varepsilon$ -almost universal, there exists at most  $\varepsilon n$  values of  $i \in [n]$  such that  $h_i(x) = h_i(y)$ , which implies that  $\Delta(H(x), H(y)) \geq (1 - \varepsilon)n$ , which proves the claim.  $\square$

## B.3 Proof of Theorem 8

*Proof.* We begin by specifying the protocol. In the pre-processing step, the client  $U$  does the following on input  $x \in [q]^k$ :

1. Generate a random  $\beta \in [n]$ .
2. Store  $(\beta, \gamma_1 = H(x_1)_\beta, \dots, \gamma_s = H(x_s)_\beta)$  and send  $x_i$  to the server  $i$  for every  $i \in [s]$ .

Server  $i$  on receiving  $x$ , saves a string  $y_i \in [q]^*$ . The server is allowed to use any computable function to obtain  $y_i$  from  $x_i$ .

During the verification phase,  $U$  does the following:

1. It sends  $\beta$  to all  $s$  servers.
2. It receives  $a_i \in [q]$  from server  $i$  for every  $i \in [s]$ . ( $a_i$  is supposed to be  $H(x_i)_\beta$ .)
3. It outputs 1 (i.e. none of the servers “cheated”) if  $a_i = \gamma_i$  for every  $i \in [s]$ , else it outputs a 0.

Similar to the one-party result, we assume that server  $i$ , on receiving the challenge, uses a computable function  $\mathcal{A}_{x,i} : [n] \times [q]^* \rightarrow [q]$  to compute  $a_i = \mathcal{A}_x(\beta, y_i)$  and sends  $a_i$  back to  $U$ .

The claim on the resource usage follows immediately from the protocol specification. Next we prove its verification guarantee. Let  $T \subseteq [s]$  be the set of colluding servers. We will prove that  $y_T$  is large by contradiction: if not, then using the list decodability of  $H$ , we will present a description of  $x_T$  of size  $< C(x_T)$ . Consider the following algorithm that uses  $y_T$  and an advice string  $v \in (\{0, 1\}^{|L|})^{|T|}$ , which is the concatenation of shorter strings  $v_i \in (\{0, 1\}^{|L|})$  for each  $i \in T$ :

1. Compute a description of  $H$  from  $n, k, \rho n, q$  and  $s$ .
2. For every  $j \in T$ , compute  $z_j = (\mathcal{A}_{x,j}(\beta, y_j))_{\beta \in [n]}$ .

3. Do the following for every  $j \in T$ : by cycling through all  $x_j \in [q]^{k/s}$ , retain the set  $\mathcal{L}_j \subseteq [q]^{k/s}$  such that for every  $u \in \mathcal{L}_j$ ,  $\Delta(H(u), z_j) \leq \rho n$ .
4. For each  $j \in T$ , let  $w_j$  be the  $v_j$ th string from  $\mathcal{L}_j$ .
5. Output the concatenation of  $\{w_j\}_{j \in T}$ .

Note that since  $H$  is  $(\rho, L)$ -list decodable, there exists an advice string  $v$  such that the algorithm above outputs  $x_T$ . Further, since  $H$  is computable, there is an algorithm  $\mathcal{E}$  that can compute a description of  $H$  from  $n, k, \rho n, q$  and  $s$ . (Note that using this description, we can generate any codeword  $H(u)$  in step 3.) Thus, we have description of  $x_T$  of size  $|y_T| + |v| + \sum_{j \in T} |\mathcal{A}_{x,j}| + |\mathcal{E}| + (s + \log(s^2 q L^s n^4) + 2 \log \log(qn) + 3)$  (where the term in parentheses is for encoding the different parameters and  $T$ ), which means that if  $|y_T| < C(x_T) - |v| - \sum_{j \in T} |\mathcal{A}_{x,j}| - |\mathcal{E}| - (s + \log(s^2 q L^s n^4) + 2 \log \log(qn) + 3) = f(x)$ , then we have a description of  $x_T$  of size  $< C(x_T)$ , which is a contradiction.  $\square$

#### B.4 Proof of Theorem 9

*Proof.* We begin by specifying the protocol. In the pre-processing step, the client  $U$  does the following on input  $x \in [q]^k$ :

1. Generate a random  $\beta \in [n]$ .
2. Store  $(\beta, \gamma = H(x)_\beta)$  and send  $x_i$  to the server  $i$  for every  $i \in [s]$ .

Server  $i$  on receiving  $x_i$ , saves a string  $y_i \in [q]^*$ . The server is allowed to use any computable function to obtain  $y_i$  from  $x_i$ . For notational convenience, we will use  $\hat{x}_i$  to denote the string  $x_i$  extended to a string in  $\mathbb{F}_q^k$  by adding zeros in positions that correspond to servers other than  $i$ .

During the verification phase,  $U$  does the following:

1. It sends  $\beta$  to all  $s$  servers.
2. It receives  $a_i \in [q]$  from server  $i$  for every  $i \in [s]$ . ( $a_i$  is supposed to be  $H(\hat{x}_i)_\beta$ .)
3. It outputs 1 (i.e. none of the servers “cheated”) if  $\gamma = \sum_{i=1}^s a_i$  else it outputs a 0.

We assume that server  $i$  on receiving the challenge, uses a computable function  $\mathcal{A}_{x,i} : [n] \times [q]^* \rightarrow [q]$  to compute  $a_i = \mathcal{A}_x(\beta, y_i)$  and sends  $a_i$  back to  $U$ . The claim on the resource usage follows immediately from the protocol specification. Next we prove its verification guarantee. Let  $T \subseteq [s]$  be the set of colluding servers. We will prove that  $y_T$  is large by contradiction: if not, then using the list decodability of  $H$ , we will present a description of  $x_T$  of size  $< C(x_T)$ .

For notational convenience, define  $\hat{x}_T = \sum_{j \in T} \hat{x}_j$  and  $\hat{x}_{\bar{T}} = \sum_{j \notin T} \hat{x}_j$ . Consider the following algorithm that uses  $y_T$  and an advice string  $v \in \{0, 1\}^{|L|}$ :

1. Compute a description of  $H$  from  $n, k, \rho, q, s$  and  $L$ .
2. Compute  $z = (\sum_{j \in T} \mathcal{A}_{x,j}(\beta, y_j))_{\beta \in [n]}$ .
3. By cycling through all  $x \in \mathbb{F}_q^k$ , retain the set  $\mathcal{L} \subseteq \mathbb{F}_q^k$  such that for every  $u \in \mathcal{L}$ ,  $\Delta(H(u), z) \leq \rho n$ .

4. Output the  $v$ th string from  $\mathcal{L}$ .

To see the correctness of the algorithm above, note that for every  $j \in [s] \setminus T$ ,  $(\mathcal{A}_{x,j}(\beta, y_j))_{\beta \in [n]} = H(\hat{x}_j)$ . Thus, if the protocol outputs 1 with probability at least  $1 - \rho$ , then  $\delta(z, H(\hat{x}_T)) \leq \rho n$ ; here we used the linearity of  $H$  to note that  $H(\hat{x}_T) = H(x) - H(\hat{x}_{\overline{T}})$ . Note that since  $H$  is  $(\rho, L)$ -list decodable, there exists an advice string  $v$  such that the algorithm above outputs  $\hat{x}_T$  (from which we can easily compute  $x_T$ ). Further, since  $H$  is computable, there is an algorithm  $\mathcal{E}$  that can compute a description of  $H$  from  $s, n, k, \rho n$  and  $q$ . Thus, we have a description of  $x_T$  of size  $|y_T| + |v| + \sum_{j \in T} |\mathcal{A}_{x,j}| + |\mathcal{E}| + (s + \log(s^2 q L n^4) + 2 \log \log(qn) + 3)$ , (where the term in parentheses is for encoding the different parameters and  $T$ ), which means that if  $|y_T| < C(x_T) - |v| - |\mathcal{A}_x| - |\mathcal{E}| - (s + \log(s^2 q L n^4) + 2 \log \log(qn) + 3) = f(x)$ , then we have a description of  $x_T$  of size  $< C(x_T)$ , which is a contradiction.  $\square$

## B.5 Proof of Theorem 10

*Proof.* We begin by specifying the protocol. As in the proof of Theorem 9, define  $\hat{x}_i$ , for  $i \in [s]$ , to be the string  $x_i$  extended to the vector in  $\mathbb{F}_q^k$ , which has zeros in the positions that do not belong to server  $i$ . Further, for any subset  $T \subseteq [s]$ , define  $\hat{x}_T = \sum_{i \in T} \hat{x}_i$ . Finally let  $RS : \mathbb{F}_q^s \rightarrow \mathbb{F}_q$  be a systematic Reed-Solomon code where  $k = 2r + e + s$ .

In the pre-processing step, the client  $U$  does the following on input  $x \in [q]^k$ :

1. Generate a random  $\beta \in [n]$ .
2. Compute the vector  $v = (H(\hat{x}_1)_\beta, \dots, H(\hat{x}_s)_\beta) \in \mathbb{F}_q^s$ .
3. Store  $(\beta, \gamma_1 = RS(v)_{s+1}, \dots, \gamma_{2r+e} = RS(v))$  and send  $x_i$  to the server  $i$  for every  $i \in [s]$ .

Server  $i$  on receiving  $x_i$ , saves a string  $y_i \in [q]^*$ . The server is allowed to use any computable function to obtain  $y_i$  from  $x_i$ .

During the verification phase,  $U$  does the following:

1. It sends  $\beta$  to all  $s$  servers.
2. For each server  $i \in [s]$ , it either receives no response or receives  $a_i \in \mathbb{F}_q$ . ( $a_i$  is supposed to be  $H(\hat{x}_i)_\beta$ .)
3. It computes the received word  $z \in \mathbb{F}_q$ , where for  $i \in [s]$ ,  $z_i = ?$  (i.e. an erasure) if the  $i$ th server does not respond else  $z_i = a_i$  and for  $s < i \leq k$ ,  $z_i = \gamma_i$ .
4. Run the decoding algorithm for  $RS$  to compute the set  $T' \subseteq [s]$  to be the error locations. (Note that by Step 2,  $U$  already knows the set  $E$  of erasures.)

We assume that server  $i$  on receiving the challenge, uses a computable function  $\mathcal{A}_{x,i} : [n] \times [q]^* \rightarrow [q]$  to compute  $a_i = \mathcal{A}_x(\beta, y_i)$  and sends  $a_i$  back to  $U$  (unless it decides not to respond).

The claim on the resource usage follows immediately from the protocol specification. We now prove the verification guarantee. Let  $T$  be the set of colluding servers. We will prove that with probability at least  $1 - \rho$ ,  $U$  using the protocol above computes  $\emptyset \neq T' \subseteq T$  (and  $|y_T|$  is large enough). Fix a  $\beta \in [n]$ . If for this  $\beta$ ,  $U$  obtains  $T' = \emptyset$ , then this implies that for every  $i \in [s]$  such that server  $i$  responds, we have  $a_i = H(\hat{x}_i)_\beta$ . This is because of our choice of  $RS$ , the decoding in Step 4 will return  $v$  (which in turn allows us to compute

exactly the set  $T' \subseteq T$  such that for every  $j \in T'$ ,  $a_j \neq H(\hat{x})_\beta$ .<sup>13</sup> Thus, if the protocol outputs a  $T' \neq \emptyset$  with probability at least  $1 - \rho$  over the random choices of  $\beta$ , then using the same argument as in the proof of Theorem 9, we note that  $\Delta(H(\hat{x}_T), (\sum_{j \in T} \mathcal{A}_{x,j}(\beta, y_j))_{\beta \in [n]}) \leq \rho n$ . Again, using the same argument as in the proof of Theorem 9 this implies that  $|y_T| \geq C(x_T) - s - \log(s^2 q L n^4) - 2 \log \log(qn) - c_0$ , for some fixed positive integer  $c_0$ .  $\square$

**Hashing Modulo a Random Prime.** We will begin with a code that corresponds to the classical Karp-Rabin hash [18]. It is known that the corresponding code  $H$  is the so called CRT code. Let  $H$  be the so called Chinese Remainder Theorem (or CRT) codes. In particular, we will consider the following special case of such codes. Let  $p_1 \leq p_2 \leq \dots \leq p_n$  be the first  $n$  primes. Consider the CRT code  $H : \prod_{i=1}^k [p_i] \rightarrow \prod_{i=1}^n [p_i]$ , where the message  $x \in \{0, 1, \dots, (\prod_{i=1}^k p_i) - 1\}$ , is mapped to the vector  $(x \bmod p_1, x \bmod p_2, \dots, x \bmod p_n) \in \prod_{i=1}^n [p_i]$ . It is known that such codes have distance  $n - k + 1$  (cf. [14]). By a simple upper bound on the prime counting function (cf. [3]), we can take  $p_n \leq 2n \log n$ . Moreover,  $\sum_{i=1}^n p_i < np_n/2$  (cf. [23]). Thus, if we pick a CRT code with  $n = k/\varepsilon^2$ , then by Theorem 4,  $H$  is  $(1 - \varepsilon, k^2(\log k - \log \varepsilon^2)/\varepsilon^4)$ -list decodable. Given any  $x \in \{0, 1, \dots, (\prod_{i=1}^k p_i) - 1\}$  and a random  $\beta \in [n]$ ,  $H(x)_\beta$  corresponds to the Karp-Rabin fingerprint (modding the input integer with a random prime). Further,  $H(x)_\beta$  can be computed in polynomial time. Thus, letting  $H$  be the CRT code in Theorem 8, we get the following:

**Corollary 17.** *For every  $\varepsilon > 0$ , there exists an  $(s, s)$ -party verification protocol with resource bound*

$$((s+1) \log k + s - \log \varepsilon^4 + s \log \log(k/\varepsilon^2), s \log k - s \log \varepsilon^2 + s + s \log \log(k/\varepsilon^2))$$

*with verification guarantee  $(\varepsilon, f)$ , where for every  $x \in \{0, 1, \dots, \prod_{i=1}^k p_i - 1\}$ ,*

$$f(x) = C(x) - c_0 (s(\log(k/\varepsilon) - \log \log(k/\varepsilon) - 1)) - c_1 \log \log \log(k/\varepsilon) - c_2$$

*for some fixed positive integers  $c_0, c_1$  and  $c_2$ . Further, all honest parties can do their computation in  $\text{poly}(n)$  time.*

**Remark 2.** *Theorem 10 can be extended to handle the case where the symbols in codewords of  $H$  are of different sizes. However, for the sake of clarity we refrain from applying CRT to the generalization of Theorem 10. Further, the results in the next subsection allow for a more efficient implementation of the computation required from the honest parties.*

**Reed-Solomon Codes.** Finally, we take  $H : \mathbb{F}_q^k \rightarrow \mathbb{F}_q^n$  to be the Reed-Solomon code, with  $n = q$ . Recall that for such a code, given message  $x = (x_0, \dots, x_{k-1}) \in \mathbb{F}_q^k$ , the codeword is given by  $H(x) = (P_x(\beta))_{\beta \in \mathbb{F}_q}$ , where  $P_x(Y) = \sum_{i=0}^{k-1} x_i Y^i$ . It is well-known that such a code  $H$  has distance  $n - k + 1$ . Thus, if we pick  $n = k/\varepsilon^2$ , then by Theorem 4,  $H$  is  $(1 - \varepsilon, 2k^2/\varepsilon^4)$ -list decodable. Let  $H$  be the Reed-Solomon code (more details in the appendix). Given any  $x \in \mathbb{F}_q^k$  and a random  $\beta \in [n]$ ,  $H(x)_\beta$  corresponds to the widely used “polynomial” hash. Further,  $H(x)_\beta$  can be computed in one pass over  $x$  with storage of only a constant number of  $\mathbb{F}_q$  elements. (Further, after reading each entry in  $x$ , by the Homer’s rule, the algorithm just needs to perform one addition and one multiplication over  $\mathbb{F}_q$ .) Thus, applying  $H$  as the Reed-Solomon code to Theorems 8 and 10 implies the following:

**Corollary 18.** *For every  $\varepsilon > 0$ ,*

---

<sup>13</sup>We will assume that  $T \cap E = \emptyset$ . If not, just replace  $T$  by  $T \setminus E$ .



- (i) *There exists an  $(s, s)$ -party verification protocol with resource bound  $((s + 1)(\log k + 2\log(1/\varepsilon) + 1), 2s(\log k + 2\log(1/\varepsilon) + 1))$  and verification guarantee  $(\varepsilon, f)$ , where for any  $x \in \mathbb{F}_q^k$ ,  $f(x) = C(x) - O(s(\log k + \log(1/\varepsilon)))$ .*
- (ii) *Assuming at most  $e$  servers do not respond to challenges, there exists an  $(r, s)$ -party verification protocol with resource bound  $((2r + e + 1)(\log k + 2\log(1/\varepsilon) + 1), 2s(\log k + 2\log(1/\varepsilon) + 1))$  and verification guarantee  $(\varepsilon, f)$ , where for any  $x \in \mathbb{F}_q^k$ ,  $f(x) = C(x) - O(s + \log k + \log(1/\varepsilon))$ .*

*Further, in both the protocols, honest parties can implement their required computation with a one pass,  $O(\log k + \log(1/\varepsilon))$  space (in bits) and  $\tilde{O}(\log k + \log(1/\varepsilon))$  update time data stream algorithm.*